

RECEIVED  
CENTRAL FAX CENTER  
JUL 19 2011

Docket No.: E0295.70199US00  
(PATENT)

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

Applicant: Stephen J. Todd et al.  
Serial No.: 10/787,337  
Confirmation No.: 3987  
Filed: February 26, 2004  
For: METHODS AND APPARATUS FOR INCREASING DATA  
STORAGE CAPACITY  
Examiner: T. S. Najee-Ullah  
Art Unit: 2453

**Certificate of Electronic Filing Under 37 CFR 1.6**  
I hereby certify that this paper (along with any paper referred to as being attached or enclosed) is being transmitted via the Office electronic filing system in accordance with § 1.8(e)(4).

Dated: July 16, 2011

Electronic Signature for Janet D'Annunzio-Ellis: /Janet D'Annunzio-Ellis/

**AMENDMENT IN RESPONSE TO NON-FINAL OFFICE ACTION**

Mail Stop Amendment  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Dear Sir:

**INTRODUCTORY COMMENTS**

In response to the Office Action dated April 15, 2011, please amend the above-identified U.S. patent application as follows:

**Amendments to the Claims** are reflected in the listing of claims which begins on page 2 of this paper.

**Remarks/Arguments** begin on page 21 of this paper.

2319661.6

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

2

Docket No.: E0295.70199US00

**AMENDMENTS TO THE CLAIMS**

Please replace all prior versions, and listings, of claims in the application with the following list of claims:

1. (Currently amended) A method of accessing a unit of data stored in a content addressable storage environment that includes a plurality of storage clusters, the method comprising acts of:  
receiving a request from a host computer to locate the unit of data previously stored in the storage environment, wherein the request identifies the unit of data via a content address that is based, at least in part, upon at least a portion of the content of the unit of data, wherein each of the plurality of storage clusters comprises a plurality of nodes that share a software utility; and  
in response to receipt of the request, determining on which one of the plurality of storage clusters the unit of data is stored based on the content address of the unit of data.
2. (Original) The method of claim 1, wherein the host computer executes an application program that stores data to and retrieves data from the storage environment, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the act of receiving is performed by the application programming interface.
3. (Original) The method of claim 1, wherein the storage environment is coupled to the host computer by at least one communication link, wherein the host computer, the storage environment, and the at least one communication link form a computer system, wherein the computer system includes an appliance that monitors access requests from the host computer, and wherein the act of receiving the request further comprises an act of receiving, at the appliance, the request from the application.
4. (Original) The method of claim 3, wherein the act of determining is performed by the appliance.

2310681.6

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

3

Docket No.: E0295.70199US00

5. (Original) The method of claim 1, wherein the host computer executes an application program that stores data on the storage environment, and wherein the act of receiving the request further comprises an act of receiving the request directly from the application program.
6. (Original) The method of claim 4, wherein the act of receiving the request further comprises an act of receiving the request at at least one of the plurality of storage clusters.
7. (Original) The method of claim 6, wherein the at least one of the plurality of storage clusters includes at least one access node that receives and processes access requests, and wherein the act of receiving the request from the application program at the at least one of the plurality of storage clusters further comprises an act of receiving the request at the at least one access node.
8. (Original) The method of claim 1, wherein the act of determining comprises an act of performing a search for the unit of data on the plurality of storage clusters.
9. (Original) The method of claim 8, wherein the act of performing a search for the unit of data further comprises an act of performing the search serially through the plurality of storage clusters until the unit of data is found.
10. (Original) The method of claim 8, wherein the act of performing a search for the unit of data further comprises an act of performing the search on each of the plurality of storage clusters in parallel.
11. (Original) The method of claim 1, wherein the act of determining is performed by at least one of the plurality of storage clusters.
12. (Original) The method of claim 1, wherein the act of determining comprises locating the unit of data on at least one of the plurality of storage clusters without performing a search.

2319561.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

4

Docket No.: E0295.70199US00

13. (Canceled)

14. (Previously presented) The method of claim 1, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the act of determining comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the time information of the content address of the unit of data.

15. (Original) The method of claim 14, wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash value of the time information of the content address of the unit of data.

16. (Previously presented) The method of claim 1, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the GUID.

17. (Original) The method of claim 16, wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash of the GUID.

18. (Previously presented) The method of claim 1, wherein the act of determining further comprises acts of:

accessing information that specifies an algorithm that was used to select on which of the plurality of storage clusters the unit of data was stored, based on the content address of the unit of data; and

applying the algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters the unit of data is stored.

2919681.6

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

5

Docket No.: E0295.70199US00

19. (Original) The method of claim 18, wherein the information specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.

20. (Original) The method of claim 19, wherein the information further specifies, for each one of the plurality of algorithms, at least one storage cluster that was in the storage environment during the period of time when the one of the plurality of algorithms was in effect.

21. (Currently amended) At least one tangible non-transitory computer readable medium encoded with instructions that, when executed on a computer system, perform a method of accessing a unit of data stored in a content addressable storage environment that includes a plurality of storage clusters, and wherein the computer system comprises the storage environment and a host computer that stores data on the storage environment, and wherein the method comprises acts of:

receiving a request from the host computer to locate the unit of data previously stored in the storage environment, wherein the request identifies the unit of data via a content address that is based, at least in part, upon at least a portion of the content of the unit of data, wherein each of the plurality of storage clusters comprises one or more nodes that share a software utility; and

in response to receipt of the request, determining on which one of the plurality of storage clusters the unit of data is stored based on the content address of the unit of data.

22. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the host computer executes an application program that stores data to and retrieves data from the storage environment, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the act of receiving is performed by the application programming interface.

23. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the computer system further comprises at least one communication link that

2319061.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

6

Docket No.: E0295.70199US00

couples the storage environment to the host computer and an appliance that monitors access requests from the host computer to the storage environment, and wherein the act of receiving the request further comprises an act of receiving, at the appliance, the request from the application.

24. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 23, wherein the act of detecting is performed by the appliance.

25. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the host computer executes an application program that stores data on the storage environment, and wherein the act of receiving the request further comprises an act of receiving the request directly from the application program.

26. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 23, wherein the act of receiving the request further comprises an act of receiving the request at at least one of the plurality of storage clusters.

27. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 26 wherein the at least one of the plurality of storage clusters includes at least one access node that receives and processes access requests, and wherein the act of receiving the request from the application program at the at least one of the plurality of storage clusters further comprises an act of receiving the request at the at least one access node.

28. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the act of determining comprises an act of performing a search for the unit of data on the plurality of storage clusters.

29. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 28, wherein the act of performing a search for the unit of data further comprises an act of performing the search serially through the plurality of storage clusters until the unit of data is found.

2319551.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

7

Docket No.: E0295.70199US00

30. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 28, wherein the act of performing a search for the unit of data further comprises an act of performing the search on each of the plurality of storage clusters in parallel.

31. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the act of determining is performed by at least one of the plurality of storage clusters.

32. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the act of determining comprises locating the unit of data on at least one of the plurality of storage clusters without performing a search.

33. (Canceled)

34. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the act of determining comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the time information of the content address of the unit of data.

35. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 34, wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash value of the time information of the content address of the unit of data.

36. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the content address includes a guaranteed unique identifier (GUID), and wherein

2310661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

8

Docket No.: E0295.70199US00

the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the GUID.

37. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 36, wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash of the GUID.

38. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 21, wherein the act of determining further comprises acts of:

accessing information that specifies an algorithm that was used to select on which of the plurality of storage clusters the unit of data was stored, based on the content address of the unit of data; and

applying the algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters the unit of data is stored.

39. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 38, wherein the information specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.

40. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 39, wherein the information further specifies, for each one of the plurality of algorithms, at least one storage cluster that was in the storage environment during the period of time when the one of the plurality of algorithms was in effect.

41. (Currently amended) An apparatus for storing data, the apparatus comprising:  
an input that receives a request from a host computer to locate a unit of data on at least one of a plurality of storage clusters in a storage environment in which the unit of data is accessible by a

2318661.6

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

9

Docket No.: E0295.70199US00

content address that is based, at least in part, upon at least a portion of the content of the unit of data, and

at least one controller coupled to the input, that:

receives the request from the input; and

in response to receipt of the request, determines on which of the plurality of storage clusters the unit of data is stored based on the content address of the unit of data,

wherein each of the plurality of storage clusters comprises one or more nodes that share a software utility.

42. (Original) The apparatus of claim 41 in combination with the storage environment, the host computer that accesses data stored in the storage environment; and a communication link that couples the host computer to the storage environment to form a computer system.

43. (Original) The apparatus of claim 42, wherein the at least one controller is disposed in the host computer.

44. (Original) The apparatus of claim 42, wherein the at least one controller is disposed in the storage environment.

45. (Original) The apparatus of claim 42 wherein the at least one controller is disposed in between the storage environment and the host computer in an appliance that monitors access requests from the host computer to the storage environment.

46. (Original) The apparatus of claim 41, wherein the host computer executes an application program that stores data to and retrieves data from the storage environment, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the at least one controller receives the request at the application programming interface.

2319001.6

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

10

Docket No.: E0295.70199US00

47. (Original) The apparatus of claim 41, wherein the host computer executes an application program that stores data in the storage environment, and wherein the at least one controller receives the request directly from the application program.

48. (Original) The apparatus of claim 44, wherein the apparatus is disposed in at least one of the plurality of storage clusters.

49. (Original) The apparatus of claim 48, wherein the at least one of the plurality of storage clusters includes at least one access node that receives and processes access requests, and wherein the at least one controller receives the request at the at least one access node.

50. (Original) The apparatus of claim 41, wherein the at least one controller performs a search for the unit of data on the plurality of storage clusters.

51. (Original) The apparatus of claim 50, wherein the at least one controller performs the search serially through the plurality of storage clusters until the unit of data is found.

52. (Original) The apparatus of claim 50, wherein the at least one controller performs the search on each of the plurality of storage clusters in parallel.

53. (Original) The apparatus of claim 41, wherein the at least one controller performs the search at at least one of the plurality of storage clusters.

54. (Original) The apparatus of claim 41, wherein the at least one controller locates the unit of data on at least one of the plurality of storage clusters without performing a search.

55. (Canceled)

2319861.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

11

Docket No.: E0295.70199US00

56. (Previously presented) The apparatus of claim 41, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the time information of the content address of the unit of data.

57. (Original) The apparatus of claim 56, wherein the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash value of the time information of the content address of the unit of data.

58. (Original) The apparatus of claim 57, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the GUID.

59. (Original) The apparatus of claim 58, wherein the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash of the GUID.

60. (Previously presented) The apparatus of claim 41, wherein the at least one controller: accesses information that specifies an algorithm that was used to select on which of the plurality of storage clusters the unit of data was stored, based on the content address of the unit of data; and

applies the algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters the unit of data is stored.

61. (Original) The apparatus of claim 60, wherein the information specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.

2319881.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

12

Docket No.: E0295.70199US00

62. (Original) The apparatus of claim 61, wherein the information further specifies, for each one of the plurality of algorithms, at least one storage cluster that was in the storage environment during the period of time when the one of the plurality of algorithms was in effect.

63. (Currently Amended) A method of accessing a unit of data stored in a content addressable storage environment that includes a plurality of storage clusters and stores data for a host computer, the method comprising acts of:

receiving a request from an application program executing on the host computer to store a unit of data, wherein each of the plurality of storage clusters comprises one or more nodes that share a software utility; and

in response to receipt of the request, selecting, based on a content address of the unit of data that is based, at least in part, upon at least a portion of the content of the unit of data, one of the plurality of storage clusters to store the unit of data in a manner that is transparent to the application program so that the application program is not made aware that the selected one of the plurality of storage clusters stores the unit of data.

64. (Original) The method of claim 63, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the act of receiving is performed by the application programming interface.

65. (Original) The method of claim 63, wherein the storage environment is coupled to the host computer by at least one communication link, wherein the host computer, the storage environment, and the at least one communication link form a computer system, wherein the computer system includes an appliance that monitors access requests from the host computer, and wherein the act of receiving the request further comprises an act of receiving, at the appliance, the request from the application.

66. (Original) The method of claim 63, wherein the act of receiving the request further comprises an act of receiving the request directly from the application program.

2319661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

13

Docket No.: E0295.70199US00

67. (Original) The method of claim 66, wherein the act of receiving the request from the application program further comprises an act of receiving the request at at least one of the plurality of storage clusters.

68. (Original) The method of claim 63, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data using a round-robin technique.

69. (Original) The method of claim 63, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on a load of at least one of the plurality of storage clusters.

70. (Original) The method of claim 63, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on an available storage capacity of each of the plurality of storage clusters.

71. (Original) The method of claim 63, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on a size of the unit of data.

72. (Canceled)

73. (Previously presented) The method of claim 63, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the act of selecting comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on the time information of the content address of the unit of data.

2318661.8

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

14

Docket No.: E0295.70199US00

74. (Original) The method of claim 73, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash value of the time information of the content address of the unit of data.

75. (Previously presented) The method of claim 63, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on the GUID.

76. (Original) The method of claim 75, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash of the GUID.

77. (Previously presented) The method of claim 63, wherein the act of selecting further comprises acts of:

applying an algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters to store the unit of data;

storing the algorithm in a record that indicates a time frame in which the algorithm was in use.

78. (Original) The method of claim 77, wherein the record specifies a plurality of algorithms used by the storage environment and at what period of time each of the plurality of algorithms was used to store units of data.

79. (Original) The method of claim 63, further comprising an act of:

storing the unit of data on the selected one of the plurality of clusters.

80. (Currently amended) At least one tangible non-transitory computer readable medium encoded with instructions that, when executed on a computer system, perform a method of

2319881.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

15

Docket No.: E0295.70199US00

accessing a unit of data stored in a storage environment in which the unit of data is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data, wherein the storage environment includes a plurality of storage clusters, each storage cluster of the plurality of storage clusters comprising one or more nodes that share a software utility, and wherein the computer system comprises the storage environment and a host computer that stores data on the storage environment, and wherein the method comprises acts of:

receiving a request from an application program executing on the host computer to store a unit of data; and

in response to receipt of the request, selecting, based on the content address of the unit of data, one of the plurality of storage clusters as to store the unit of data in a manner that is transparent to the application program so that the application program is not made aware that the selected one of the plurality of storage clusters stores the unit of data.

81. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the act of receiving is performed by the application programming interface.

82. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the computer system further comprises at least one communication link that couples the storage environment to the host computer to the storage environment and an appliance that monitors access requests from the host computer, and wherein the act of receiving the request further comprises an act of receiving, at the appliance, the request from the application.

83. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the act of receiving the request further comprises an act of receiving the request directly from the application program.

2319551.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

16

Docket No.: E0295.70199US00

84. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 83, wherein the act of receiving the request from the application program further comprises an act of receiving the request at at least one of the plurality of storage clusters.

85. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data using a round-robin technique.

86. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on an available storage capacity of each of the plurality of storage clusters.

87. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on a load of at least one of the plurality of storage clusters.

88. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on a size of the unit of data.

89. (Canceled)

90. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the act of selecting comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on the time information of the content address of the unit of data.

2319661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

17

Docket No.: E0295.70199US00

91. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 90, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash value of the time information of the content address of the unit of data.

92. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on the GUID.

93. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 91, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash of the GUID.

94. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the act of selecting further comprises acts of:

applying an algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters to store the unit of data;

storing the algorithm in a record that indicates a time frame in which the algorithm was in use.

95. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 94, wherein the record specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.

96. (Currently amended) The at least one tangible non-transitory computer readable medium of claim 80, wherein the method further comprises an act of:

storing the unit of data on the selected one of the plurality of clusters.

2319601.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

18

Docket No.: E0295.70199US00

97. (Currently amended) An apparatus for storing data, the apparatus comprising:

an input that receives a request from an application program executing on a host computer to store a unit of data on at least one of a plurality of storage clusters in a storage environment in which the unit of data is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data, and

at least one controller, coupled to the input, that:

receives the request from the input; and

in response to receipt of the request, selects, based on the content address of the unit of data, one of the plurality of storage clusters to store the unit of data in a manner that is transparent to the application program so that the application program is not made aware that the selected one of the plurality of storage clusters stores the unit of data,

wherein each of the plurality of storage clusters comprises one or more nodes that share a software utility.

98. (Original) The apparatus of claim 97 in combination with the storage environment, the host computer that accesses data stored in the storage environment; and a communication link that couples the host computer to the storage environment to form a computer system.

99. (Original) The apparatus of claim 98, wherein the at least one controller is disposed in the host computer.

100. (Original) The apparatus of claim 98, wherein the at least one controller is disposed in the storage environment.

101. (Original) The apparatus of claim 98 wherein the at least one controller is disposed in between the storage environment and the host computer in an appliance that monitors access requests from the host computer to the storage environment.

2318661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

19

Docket No.: E0295.70199US00

102. (Original) The apparatus of claim 99, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the at least one controller receives the request at the application programming interface.

103. (Original) The apparatus of claim 97, wherein the at least one controller receives the request directly from the application program through the input.

104. (Original) The apparatus of claim 100, wherein the at least one controller receives the request at at least one of the plurality of storage clusters.

105. (Original) The apparatus of claim 97, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data using a round-robin technique.

106. (Original) The apparatus of claim 97, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on an available storage capacity of each of the plurality of storage clusters.

107. (Original) The apparatus of claim 97, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based on a load of at least one of the plurality of storage clusters.

108. (Original) The apparatus of claim 97, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based on a size of the unit of data.

109. (Canceled)

110. (Previously presented) The apparatus of claim 97, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the at least

2319661.6

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

20

Docket No.: E0295.70199US00

one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on the time information of the content address of the unit of data.

111. (Original) The apparatus of claim 110, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash value of the time information of the content address of the unit of data.

112. (Previously presented) The apparatus of claim 97, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on the GUID.

113. (Original) The apparatus of claim 112, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash of the GUID.

114. (Previously presented) The apparatus of claim 97, wherein the at least one controller: applies an algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters to store the unit of data; and stores the algorithm in a record that indicates a time frame in which the algorithm was in use.

115. (Original) The apparatus of claim 114, wherein the record specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.

116. (Original) The apparatus of claim 97, wherein the at least one controller stores the unit of data on the selected one of the plurality of clusters.

2919661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

21

Docket No.: E0295.70199US00

### REMARKS

In response to the Office Action dated April 15, 2011, Applicant respectfully requests reconsideration. Claims 1-12, 14-32, 34-54, 56-71, 73-88, 90-108, and 110-116 were previously pending in this application. By this amendment, Applicant is canceling no claims and adding no claims. Claims 1, 21, 41, 63, 80 and 97 are amended. As a result, claims 1-12, 14-32, 34-54, 56-71, 73-88, 90-108, and 110-116 remain pending for examination with claims 1, 21, 41, 63, 80 and 97 being independent. No new matter has been added.

#### Rejections Under 35 U.S.C. §101

The Office Action rejects claims 21-32, 34-40, 80-88 and 90-96 under 35 U.S.C. §101 because the claimed invention is purportedly directed to non-statutory subject matter. Applicant has amended claims 21-32, 34-40, 80-88 and 90-96 for clarity and respectfully asserts that these amendments render moot the rejection of claims 21-32, 34-40, 80-88 and 90-96 under 35 U.S.C. § 101. Accordingly, reconsideration and withdrawal of these rejections is respectfully requested.

#### Rejections Under 35 U.S.C. §103

The Office Action rejects claims 1-12, 14-32, 34-54, 56-67, 69-71, 73-84, 86-88, 90-104, 106-109, and 110-116 under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 6,173,374 (Heil) in view of U.S. Patent Publication No. 2005/0005066 (Nakayama). These rejections are respectfully traversed.

##### I. Independent Claim 1

The rejection of independent claim 1 should be withdrawn because Heil and Nakayama, alone or in combination, fail to teach or suggest every limitation of claim 1. Specifically, the cited references do not disclose “identif[ying] the unit of data via a content address that is based, at least in part, upon at least a portion of the content of the unit of data.”

Moreover, without acceding to the propriety of the rejections, independent claim 1 is herein amended to clarify the claimed subject matter. Each independent claim of the present application, as amended, recites, “each of the plurality of storage clusters comprises *a plurality of nodes that*

2319661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

22

Docket No.: E0295.70199US00

*share a software utility.*" Support for this amendment may be found throughout the original disclosure, for example, at page 8, lines 8-24. The cited references fail to disclose this additional limitation.

A. *The cited references, alone or in combination, fail to disclose "identif[ying] the unit of data via a content address that is based, at least in part, upon at least a portion of the content of the unit of data"*

Neither Heil, nor Nakayama, alone or in combination, teach or suggest "identif[ying] the unit of data via a content address that is based, at least in part, upon at least a portion of the content of the unit of data," as recited by claim 1. The Office Action (page 6) concedes that Heil fails to disclose this limitation, but asserts that Nakayama teaches it, citing ¶85, ¶¶22-26, ¶¶30-33 and ¶¶45-47. Applicant respectfully disagrees.

The cited portions of Nakayama are directed to a "guarantee code" that is "*based on* an attribute (for example, *logical address* and error-correcting redundant code) of the data which is requested to be written" (¶23). Nakayama is thus basing the guarantee code on the address of the data being written. Nakayama does not teach the guarantee code being used as a "content address." Nor would it be possible for the guarantee code to be used as a "content address" of the data because the address must already exist in order to generate the guarantee code. The guarantee code is merely a means to ensure that data is copied accurately (¶22) and not a "content address." The only address discussed by Nakayama is the "logical address," which is not "based, at least in part, upon at least a portion of the content of the unit of data." Therefore, Heil and Nakayama, alone or in combination, fail to disclose at least this limitation of claim 1.

B. *The cited references, alone or in combination, fail to disclose "each of the plurality of storage clusters comprises a plurality of nodes that share a software utility"*

Neither Heil nor Nakayama, alone or in combination, teach or suggest "each of the plurality of storage clusters comprises a plurality of nodes that share a software utility," as recited by amended claim 1. The Office Action (page 6) asserts that the storage cluster of Heil maps to the "plurality of storage clusters" of claim 1 and that each individual remote disk of Heil corresponds to

2319661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

23

Docket No.: E0295.70199US00

"each of the plurality of storage clusters." However, each remote disk is associated with a single node (see, e.g., Heil FIG. 1). Therefore, Heil does not disclose "each of the plurality of storage clusters comprises *a plurality of nodes that share a software utility*." Nakayama cannot cure this deficiency because the reference discloses a storage system capable of remote copying between different storage systems (¶2). Each of the storage systems is associated with a single node (see FIG. 1). Thus, Nakayama does not disclose "each of the plurality of storage clusters comprises *a plurality of nodes that share a software utility*." Therefore, Heil and Nakayama, alone or in combination, fail to disclose at least this limitation of claim 1.

For at least the foregoing reasons, Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 1. Accordingly, withdrawal of the rejection of independent claim 1 is respectfully requested.

Claims 2-12 and 14-20 depend from claim 1 and are allowable for at least the same reasons.

#### Independent Claim 21

The rejection of independent claim 21 should be withdrawn because Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 21. Specifically, the limitations "each of the plurality of storage clusters comprises one or more nodes that share a software utility" and "identifies the unit of data via a content address that is based, at least in part, upon at least a portion of the content of the unit of data" are not disclosed by the cited references for reasons that should be clear from the above discussion of the references in connection with of claim 1.

For at least the foregoing reasons, Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 1. Accordingly, withdrawal of the rejection of independent claim 21 is respectfully requested.

Claims 22-32 and 34-40 depend from claim 21 and are allowable for at least the same reasons.

#### Independent Claim 41

The rejection of independent claim 41 should be withdrawn because Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 41. Specifically, the limitations

2310661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

24

Docket No.: E0295.70199US00

"each of the plurality of storage clusters comprises one or more nodes that share a software utility" and "the unit of data is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data" are not disclosed by the cited references for reasons that should be clear from the above discussion of the references in connection with of claim 1.

For at least the foregoing reasons, Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 41. Accordingly, withdrawal of the rejection of independent claim 41 is respectfully requested.

Claims 42-62 depend from claim 41 and are allowable for at least the same reasons.

Independent Claim 63

The rejection of independent claim 63 should be withdrawn because Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 63. Specifically, the limitations "each of the plurality of storage clusters comprises one or more nodes that share a software utility" and "a content address of the unit of data that is based, at least in part, upon at least a portion of the content of the unit of data" are not disclosed by the cited references for reasons that should be clear from the above discussion of the references in connection with of claim 1.

For at least the foregoing reasons, Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 63. Accordingly, withdrawal of the rejection of independent claim 63 is respectfully requested.

Claims 64-71 and 73-79 depend from claim 63 and are allowable for at least the same reasons.

Independent Claim 80

The rejection of independent claim 80 should be withdrawn because Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 80. Specifically, the limitations "each of the plurality of storage clusters comprises one or more nodes that share a software utility" and "the unit of data is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data" are not disclosed by the cited references for reasons that should be clear from the above discussion of the references in connection with of claim 1.

2319661.5

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

25

Docket No.: E0295.70199US00

For at least the foregoing reasons, Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 80. Accordingly, withdrawal of the rejection of independent claim 80 is respectfully requested.

Claims 81-88 and 90-96 depend from claim 80 and are allowable for at least the same reasons.

Independent Claim 97

The rejection of independent claim 97 should be withdrawn because Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 97. Specifically, the limitations "each of the plurality of storage clusters comprises one or more nodes that share a software utility" and "the unit of data is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data" are not disclosed by the cited references for reasons that should be clear from the above discussion of the references in connection with of claim 1.

For at least the foregoing reasons, Heil and Nakayama, alone or in combination, fail to disclose every limitation of claim 97. Accordingly, withdrawal of the rejection of independent claim 97 is respectfully requested.

Claims 98-108 and 110-116 depend from claim 97 and are allowable for at least the same reasons.

General Comments on Dependent Claims

Each of the dependent claims depends from a base claim that is believed to be in condition for allowance, and Applicants believe that it is unnecessary at this time to argue the allowability of each of the dependent claims individually. Applicants do not, however, necessarily concur with the interpretation of the dependent claims as set forth in the Office Action, nor do Applicants concur that the basis for the rejection of any of the dependent claims is proper. Therefore, Applicants reserve the right to specifically address the patentability of the dependent claims in the future, if deemed necessary.

2319661.6

Application No. 10/787,337  
Reply to Office Action of April 15, 2011

26

Docket No.: E0295.70199US00

**CONCLUSION**

A Notice of Allowance is respectfully requested. The Examiner is requested to call the undersigned at the telephone number listed below if this communication does not place the case in condition for allowance.

If this response is not considered timely filed and if a request for an extension of time is otherwise absent, Applicant hereby requests any necessary extension of time. If there is a fee occasioned by this response, including an extension fee, the Director is hereby authorized to charge any deficiency or credit any overpayment in the fees filed, asserted to be filed or which should have been filed herewith to our Deposit Account No. 23/2825, under Docket No. E0295.70199US00 from which the undersigned is authorized to draw.

Dated: July 15, 2011

Respectfully submitted,

By /Edmund J. Walsh/  
Edmund J. Walsh  
Registration No.: 32,950  
WOLF, GREENFIELD & SACKS, P.C.  
600 Atlantic Avenue  
Boston, Massachusetts 02210-2206  
617.646.8000

2319601.6



# Wolf Greenfield

RECEIVED  
CENTRAL FAX CENTRE  
JUL 19 2011

**F A X C O V E R**

\*\*\*\*\* OFFICIAL FAX \*\*\*\*\*

**Date:** July 19, 2011 **Number of pages (including cover):** 29

**To:** US Patent and Trademark Office  
**Attn:** Examiner T.S. Nafee-Ullah

**Fax No.:** ✆ 571.273.8300 (Central)  
And ✆ 571.270-6013 (Examiner Nasee-Ullah)

Serial No.: 10/787,337-Conf. #3987

## **METHODS AND APPARATUS FOR INCREASING DATA STORAGE CAPACITY**

**From:** Edmund J. Walsh

**Direct dial:** 617.646.8000

**Our File #:** E0295.70199US00

**CERTIFICATE OF FACSIMILE TRANSMISSION 37 C.F.R. 51.8(a)**

The undersigned hereby certifies that this document is being transmitted via facsimile to the attention of US Patent and Trademark Office, FAX number 571.273.8300, and FAX number 571.270.6013 at the United States Patent and Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450, in accordance with 37 C.F.R. 51.6(d), on July 19, 2011.

/Danielle Calder/  
Danielle A. Calder

**ORIGINAL DOCUMENTS WILL NOT BE MAILED.**

**MESSAGE:** Transmitted herewith is/are . (1) Applicant Initiated Interview Request Form (1 page)  
(2) Interview Agenda (1 page)  
(3) Copy of Amendment filed on 07/15/11 (26 pages)

This transmission contains confidential information intended for use only by the above-named recipient. Reading, discussing, distributing, or copying this message by anyone other than the named recipient, or his or her employees or agents, is strictly prohibited. If you have received this fax in error, please notify us immediately by telephone (collect), and return the original message to us at the address below via the U.S. Postal Service.

**IF YOU DID NOT RECEIVE ALL OF THE PAGES OF THIS TRANSMISSION, OR IF ANY OF THE PAGES ARE ILLEGIBLE, PLEASE CALL 617.646.8000 IMMEDIATELY.**

Wolf Greenfield Fax Number: 617.646.8646

Wolf, Greenfield & Sacks, P.C. | 600 Atlantic Avenue | Boston, Massachusetts 02210-2206  
617.646.8000 | fax 617.646.8646 | [www.wolfgreenfield.com](http://www.wolfgreenfield.com)

PATENTS TRADEMARKS COPYRIGHTS TECHNOLOGY TRANSFERS LITIGATION